Abnormal

RESEARCH REPORT

2026 State of Behavioral Security

The Challenge of Defining Normal Across People, Machine Identities, and Autonomous AI Agents



Behavioral Security Is Mainstream; the Hard Part Comes Next

Behavioral security has become a core security discipline.

But broad adoption has exposed a more consequential challenge: can organizations define normal precisely enough to separate meaningful risk from noise, turn anomalies into confident action, and extend that understanding to a rapidly expanding population of non-human identities and autonomous AI agents?

The 2026 State of Behavioral Security examines where the market has matured, where operational friction remains, and where security teams are placing their bets next. The findings point to a clear shift: with mainstream adoption, **behavioral maturity has become the defining problem to solve.**

Key Findings at a Glance

Behavioral Security Goes Mainstream

81.1% of respondents rely on behavioral anomaly detection to a major or extensive degree, while **87.4%** say behavioral baselining is mostly or comprehensively operationalized.

The Knowing-Normal Gap

47.8% cite false positives and **43.6% cite difficulty defining normal** as the **two leading barriers** to broader behavioral baselining implementation.

The Understand-to-Act Gap

90.1% say understanding **why an anomaly triggered is very or critically important**, but **51.2%** typically need **at least four hours** to fully investigate and remediate one.

The Machine Identity Confidence Paradox

71.6% say they are better prepared to detect anomalous behavior from non-human identities than human identities, even as behavioral-baseline coverage and effectiveness remain uneven across NHIs.

The AI Agent Investment Paradox

43% are putting their highest behavioral-monitoring **investment into AI agents**—more than 3× the next category—**yet 46.3%** say they are **not at all effective at maintaining AI agent baselines.**

Methodology

525 cybersecurity professionals

Organizations of **1,000+** employees

6 countries

85% CISO-level or SOC leadership

4 in-depth CISO interviews

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Behavioral Detection Is a Core Security Discipline

Adoption Hasn't Solved the Hardest Problem: Knowing Normal

Detection Is Only the Beginning

Behavioral Security Is Expanding Beyond Human Identities

AI Agents Are the Next Behavioral-Security Frontier

Conclusion: What Comes After Behavioral Adoption

Behavioral Detection Is a Core Security Discipline

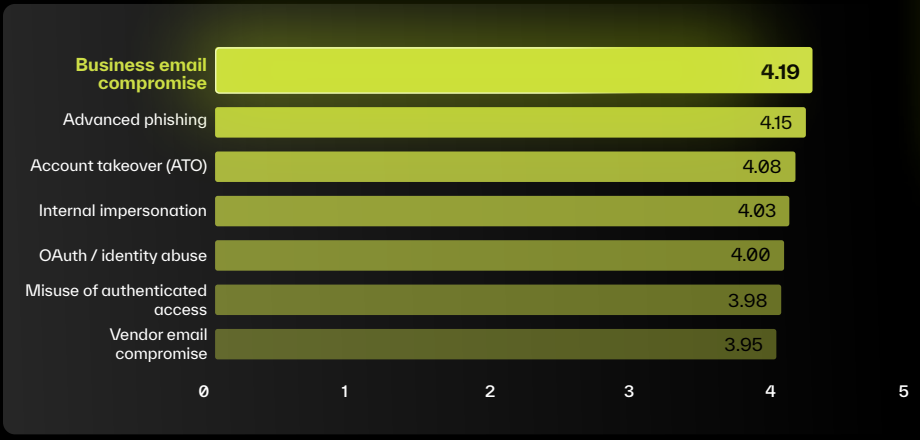
Behavioral security is firmly established as part of the modern security stack.

More than eight in ten organizations (81.1%) rely on behavioral anomaly detection to a major or extensive degree, while **87.4% say behavioral baselining is mostly or comprehensively operationalized**.

91.2% say behavioral baselining will be very or critically important two years from now.

Effectiveness of Organizations at Detecting Each Type of Threat

On a scale of 1 (least) to 5 (most).



That practitioner described behavioral techniques already embedded across Microsoft, identity, authentication, and email-security environments as increasingly routine components of how security technologies assess risk.

A second CISO described a similar pattern around identity, where the organization already uses "several layers of behavioral detection," from capabilities built into identity providers to broader network and cloud-perimeter analysis.

Organizations also continue to rely on known malicious indicators, reputation, rules, signatures, and threat intelligence. The importance of behavioral baselining becomes especially clear when a threat does **not** contain the signals conventional detection is designed to recognize.

When asked which approach primarily enables detection of low-signal, suspicious activity without known malicious indicators, **45.1% of respondents selected behavioral anomaly detection**; nearly as many as threat intelligence and signature- or rule-based detection combined.

According to one CISO at a business advisory and financial services organization:

> **Behavioral detections in general are kind of table stakes in a lot of tools."**

45%
Behavioral anomaly detection

25%
Threat intelligence

22%
Signature- or rule-based detection

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Taken together, the findings suggest an important shift in the enterprise detection model. **Known malicious indicators remain important, but security teams increasingly need to detect suspicious activity when those indicators are absent. Behavioral detection has become the leading approach for doing so.**

That matters because the absence of a known-bad signal does not mean the absence of risk. A legitimate credential can still be misused. An authenticated identity can still behave in an unexpected way. Activity can be technically permitted while remaining inconsistent with what that user, application, or system normally does.

Behavior provides another way to evaluate those situations: not only by asking whether an action matches something already known to be malicious, but whether it makes sense in context; or, **is it normal?**

How well can organizations define what normal actually looks like?

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Adoption Hasn't Solved the Hardest Problem: Knowing Normal

Most organizations have already embraced behavioral baselining. The harder problem is establishing a baseline of **normal behavior** that is precise enough to distinguish meaningful deviation from the ordinary variability of enterprise activity.

The Knowing-Normal Gap

Yet the two most commonly cited barriers to expanding it are:

48% — Too many false positives
44% — Difficulty defining "normal"

For human identities especially, "normal" rarely means one fixed pattern.

For that organization, the problem is concrete. Its identity population includes employees as well as physicians and other authorized users whose activity can vary dramatically. Someone may access the environment several times in three months, then disappear for six months or a year. A rule that interprets prolonged inactivity as abnormal could therefore flag behavior that is entirely legitimate.

That variability is what turns baselining from a deployment problem into a context problem. The more normal activity changes by person, role, application, location, or circumstance, the harder it becomes to draw a useful boundary between unusual and risky.

87%
of organizations say behavioral baselining is **mostly or comprehensively operationalized.**

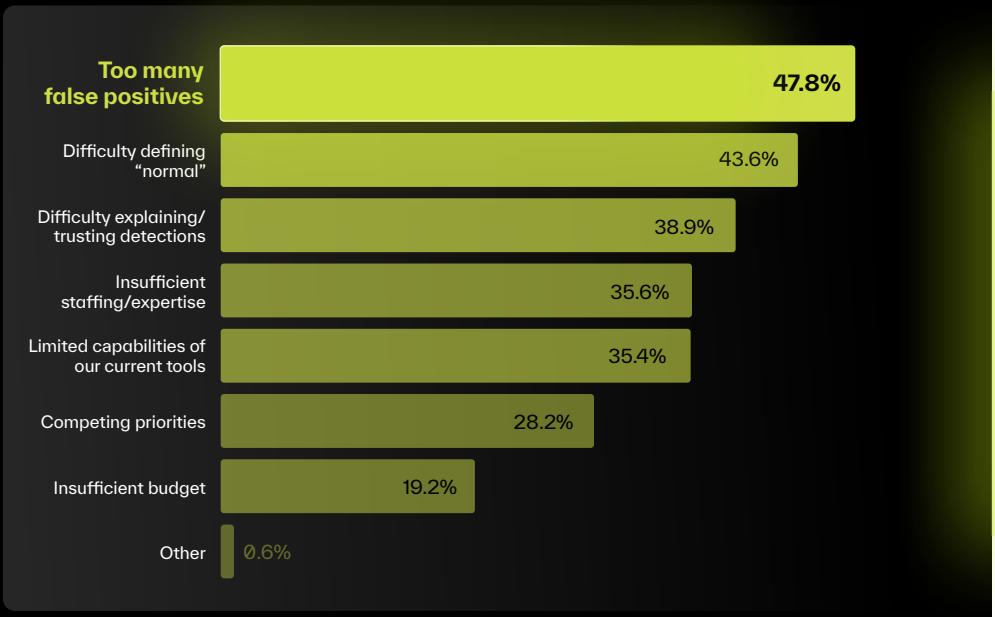
> There's not a consistent flow of behavior when it comes to humans, that we can say, this is what normal looks like all the time."
>
> — CISO, healthcare organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

What gets in the way of broader behavioral baselining?

Greatest Obstacles to Implementing Behavioral Baselining

Percent of respondents who selected each option.



> Too many false positives start to erode confidence."
>
> — CISO, healthcare organization

The three leading barriers center on the quality and usability of behavioral baselining:

- Reducing false positives
- Defining normal
- Explaining or trusting detections

There Is No Single Formula for Normal

The data also shows that organizations do not rely on one universal method to establish a baseline.

Automatically learned baselines are common in several mature environments—used by **53.0% for SaaS, 51.6% for identity and authorization, and 48.2% for email**—while predefined rules continue to play a substantial role alongside them. Organizations frequently use both approaches rather than treating behavioral learning as a replacement for established logic.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

How Organizations Establish Normal

The mix makes sense. Some behaviors can be bounded with explicit policies. Others require systems to learn what is typical over time. In practice, mature behavioral security increasingly depends on combining those approaches with enough identity and environmental context to determine whether a deviation is merely unusual or actually suspicious.

That distinction becomes especially important when an attacker does not need to break through a control at all.

> **You're not being hacked. They're logging in."**
>
> — CISO, healthcare organization

The same CISO described attackers using legitimate credentials and then appearing normal once inside the environment: "They're sticking with the credentials and they're logging in, and then they look normal."

When valid access and malicious access can look superficially similar, detecting deviation is only useful if the baseline contains enough context to tell the difference.

The behavioral maturity challenge is knowing normal well enough to separate meaningful risk from ordinary variation.

SaaS

53% automatically learned

Identity & authorization

52% automatically learned

Email

48% automatically learned

Machine identities

36% automatically learned

53% do not establish normal at all

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Detection Is Only the Beginning

Finding an anomaly is only the first step. The harder operational test is whether security teams can understand what happened, determine whether it represents real risk, and respond quickly enough to matter.

On the surface, trust is relatively high: **65.3% of respondents say analysts trust behavior-based detections enough to act** without additional manual validation always or most of the time.

But unconditional trust is rare. Only **7.8% say analysts can always act without additional validation.**

The Understand-to-Act Gap

Detect	Anomaly identified
Understand	90.1% say knowing why an anomaly triggered is very or critically important
Trust	**7.8% can always act without additional manual validation**
Act + Resolve	51.2% typically need 4+ hours to fully investigate and remediate a behavioral anomaly

The importance placed on explanation is revealing. Nine in ten respondents say understanding why an anomaly was triggered is very or critically important when determining how to respond. Separately, **38.9% identify difficulty explaining or trusting detections as an obstacle to expanding behavioral baselining.**

The CISO interviews show why. An anomaly can establish that something changed without establishing what that change means.

> All we know is it's misbehaving."
>
> — CISO, healthcare organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

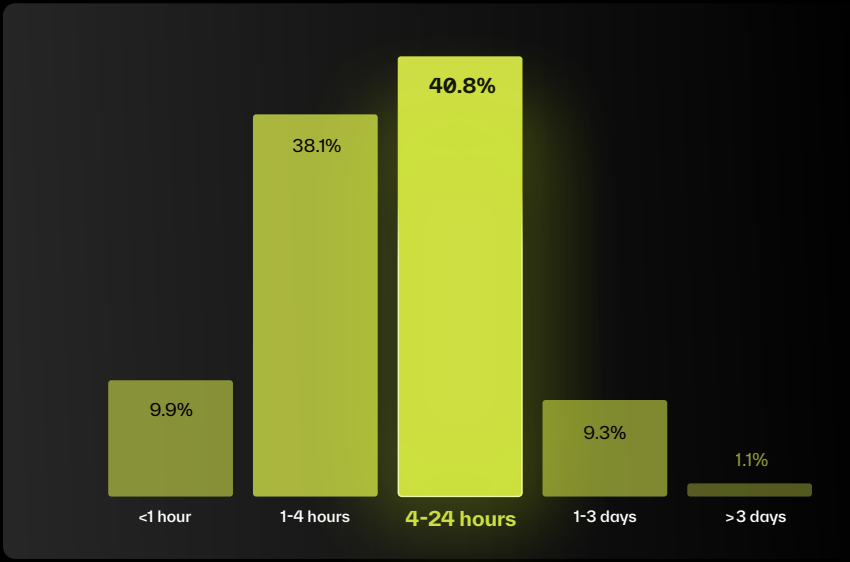
In the example mentioned by the CISO, a device could be automatically isolated because its behavior crossed a threshold. But the security team still had to identify the owner, understand what the device was used for, and determine what should happen next.

In other environments, making that judgment can require knowledge well beyond the SOC. One healthcare security leader described behavioral anomalies being reviewed by privacy officers, HR teams, and clinicians.

Investigation Still Takes Time

Typical Time to Investigate/Remediate Behavioral Detection

Length of time elapsed before investigation or remediation finished.



More than half of organizations—51.2%—typically need at least four hours to move from initial detection through full investigation and remediation.

Over 10% of organizations need at least one full day to do the same.

That time reflects more than technical investigation. Behavioral detections frequently require context: who or what generated the activity, what that identity normally does, whether permissions or responsibilities recently changed, and whether an unusual action has a legitimate explanation.

The result can be a process that remains heavily dependent on human judgment and information sharing.

> We have privacy officers, we have HR, we have clinicians look at that and decide if that's normal behavior or not, and then take action."
>
> — CISO, healthcare organization

> This is how we learn... by sharing information... But that's... a very manual, serial process."
>
> — CISO, healthcare organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

At enterprise scale, that model becomes difficult to sustain.

Another CISO described the volume of security data as “phenomenal,” arguing that it can no longer be handled without some degree of automation and AI-assisted detection.

Organizations appear to recognize the bottleneck. Improved automation is the most commonly cited priority for strengthening behavioral baselining, selected by 53.9% of respondents, while 37.1% prioritize faster investigations.

54%
prioritize improved automation

37%
prioritize faster investigations

Visibility Is Not Understanding

79% of respondents say they can easily produce a unified view of anomalous activity across major communication and identity environments.

That is a meaningful sign of progress. But aggregating anomalies is not the same as explaining them. A unified view can show what changed. Effective response still depends on understanding why it changed, whether it matters, and what action is justified.

The maturity gap is no longer simply between seeing and missing an anomaly. It is between detecting unusual behavior and understanding it well enough to act with confidence and without delay.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Behavioral Security Is Expanding Beyond Human Identities

Behavioral security may be mainstream, but its reach is uneven.

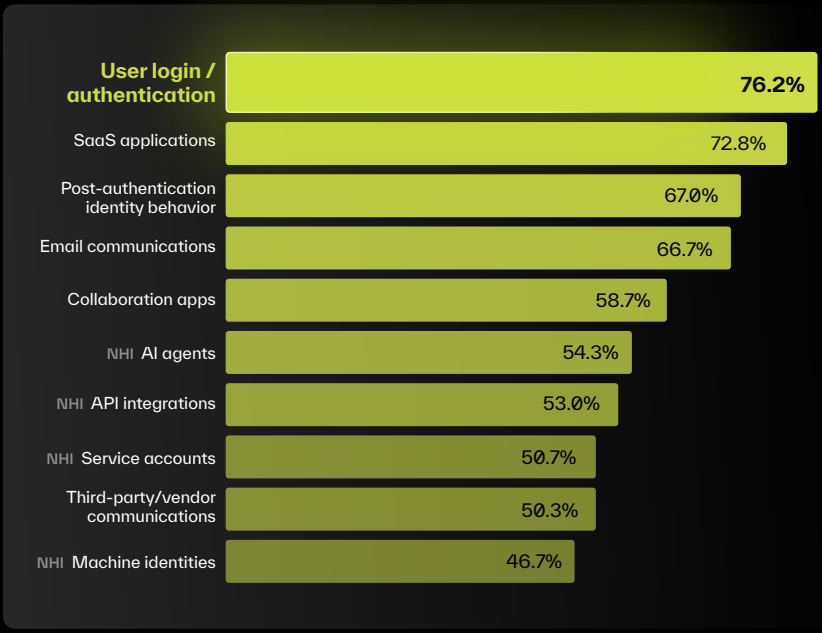
Organizations are most likely to maintain behavioral baselines around areas tied closely to human activity: **76.2% for user login and authentication, 67.0% for post-authentication identity behavior, and 66.7% for email communications.** Coverage is also high for **SaaS applications**, at **72.8%.**

Across most non-human categories, it drops closer to half.

Where Organizations Maintain Behavioral Baselines Today

Behavioral Baseline Monitoring Coverage Across Systems and Activities

Percent of respondents who maintain baselines in each area.



Behavioral baselining is most established around users and mature application environments, while non-human identities (NHI) are concentrated toward the lower end of baseline coverage.

> The non-human identities... are really where the new playground of risk is surfacing."
>
> — CISO, B2B software organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

That difference matters because non-human activity is not peripheral to modern organizations. Applications, workloads, APIs, service accounts, and other machine identities routinely access systems and data without a person directly performing each action.

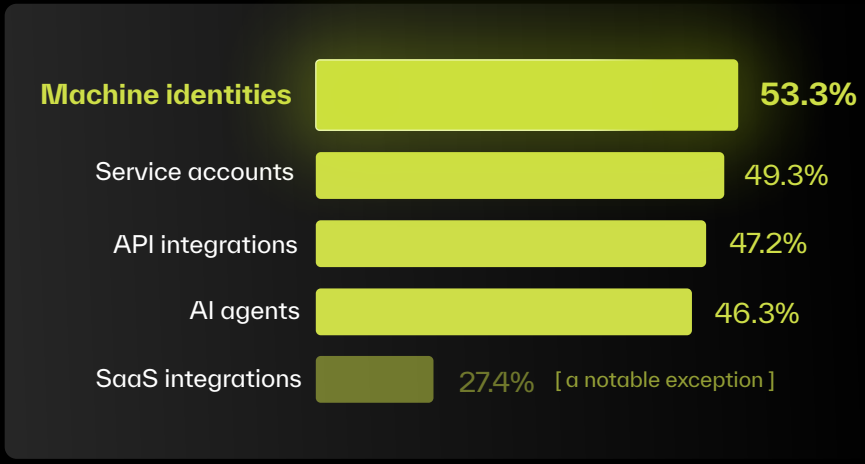
The Machine Identity Confidence Paradox

Interestingly, the coverage gap does not translate into lower perceived preparedness.

71.6% of respondents say their organization is somewhat or much better prepared to detect anomalous behavior from non-human identities than from human identities. Yet behavioral-baseline coverage remains incomplete across machine identities, service accounts, APIs, and AI agents. Effectiveness is similarly uneven:

Organizations Saying They Are Not at All Effective at Establishing and Maintaining Behavioral Baselines for:

Percent of respondents who selected each option.



The findings do not necessarily mean security teams are consistently mistaken about their preparedness. One possible explanation is that many traditional non-human identities still perform narrower, more predictable functions than people do. A service account with a defined purpose and set of permissions may make deviation easier to recognize even if behavioral monitoring around it is less comprehensive. But AI agents are bound to take on far more advanced functions.

That distinction becomes important as behavioral security expands further. Traditional machine activity may often be constrained by relatively defined functions. **AI agents can introduce a different degree of autonomy, and a different problem for defining normal.**

The interviews provide some support for that interpretation.

> **We know what a service account's supposed to do and what it's not supposed to do. The challenge with [AI] agents is: we don't know what we don't know."**
>
> — CISO, healthcare organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

How We Defined Machine Identities and AI Agents

Machine identities are non-human credentials or principals used by workloads, devices, services, or applications to authenticate to and access enterprise resources.

AI agents are goal-directed software actors that can make decisions about how to complete a task and take actions with some degree of autonomy.

An AI agent may use a machine identity to authenticate, so the distinction is not "machine versus AI." It is identity versus actor: the identity used to obtain access versus the software actor deciding what to do with that access.

One CISO offered a concrete example of conventional machine behavior: access could require not only valid credentials but an authenticated machine carrying the organization's certificate. If that machine was rebuilt, moved, or taken off the network, its behavior had changed and a different policy could apply.

The SaaS Advantage

SaaS stands out across two related but distinct measures. SaaS applications rank near the top for baseline coverage and use of automatically learned baselines, while SaaS integrations report substantially stronger baseline effectiveness than the other non-human categories measured.

One CISO described a relatively controlled environment in which enterprise applications and OAuth integrations are audited, access routes through established identity systems, and users cannot simply approve arbitrary OAuth grants.

That does not prove why SaaS performs better in the survey. But the interviews suggest one possible explanation: **SaaS environments benefit from years of investment in identity, access, and application governance that make activity easier to attribute and constrain.** Established controls around SSO, provisioning, OAuth permissions, and application approval can give security teams a clearer picture of who or what should have access, and under what conditions.

73%
maintain behavioral baselines for SaaS applications.

53%
use automatically learned baselines for SaaS—the highest level measured.

62%
rate themselves very or extremely effective at establishing and maintaining baselines for SaaS integrations.

Behavioral security is expanding beyond human identities. But the further it moves into non-human activity, the more important the nature of the actor becomes. Traditional machine identities are only part of that story.

The newest challenge is more autonomous and unpredictable.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

AI Agents Are the Next Behavioral-Security Frontier

Security teams are making their priorities clear.

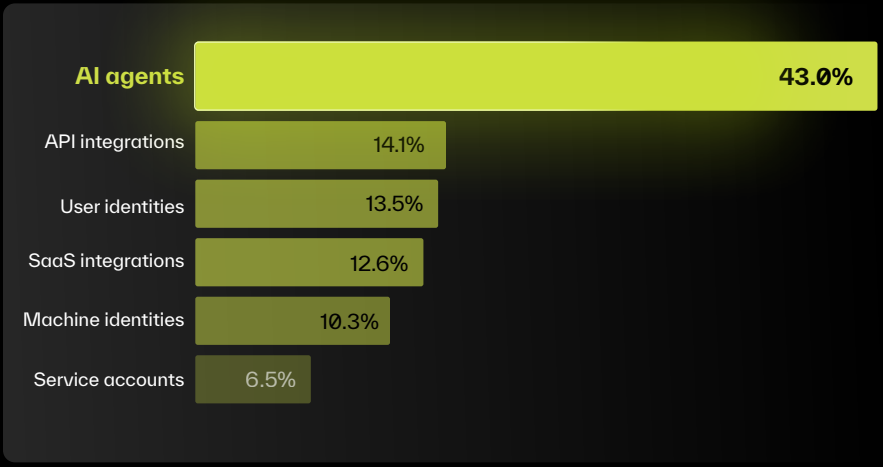
When asked to identify the **single area that will receive the greatest behavioral-monitoring investment over the next 24 months, 43.0% chose AI agents.** The next-highest category—API integrations—received just 14.1%.

That is not a narrow lead. **AI agents received more than three times the response of any other category.**

The AI Agent Investment Paradox

Where will organizations make their greatest behavioral-monitoring investment?

Percent of respondents who selected each option.



AI agents have separated decisively from every other category as the next behavioral-monitoring investment priority.

The market is not starting from zero. **54.3% already say they maintain behavioral baselines around AI agents,** and 30.9% name AI agent coverage as one of their top three priorities for improving behavioral baselining.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Current effectiveness, however, is sharply divided. While 41.9% rate their organizations as very or extremely effective at establishing and maintaining AI agent baselines, **46.3% say they are not at all effective.**

The tension is clear: investment grows rapidly while organizations are still working out what effective monitoring should look like.

One CISO described essentially that dynamic among security leaders:

> **It is the number one planned investment for practically every CISO I've talked to... for next year."**

Yet the same CISO added:

> **There is no plan. Like, we don't know what the solutions are, but we're all gonna put resources to go after it."**

— CISO, B2B software organization

Autonomy Changes the Behavioral Question

The challenge is not simply extending existing machine-identity controls to one more category.

AI agents can act on behalf of people, access information across multiple systems, make decisions about how to complete assigned tasks, and generate additional actions that were not individually specified by a human.

That complicates a security model built around relatively stable relationships between identity, access, and action. An agent may operate under delegated permissions, retrieve information from systems its human principal does not access directly, or take a series of intermediate actions in pursuit of a legitimate goal.

43%
say AI agents are the **#1 area for future investment**

46%
say they are **not at all effective** at maintaining AI agent baselines

As one CISO described it, agents are already beginning to blur conventional attribution:

> **It's blurring the lines between what is truly the agent and what is the human, because it's acting on behalf of people."**

— CISO, B2B software organization

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

The security problem therefore becomes more nuanced than determining whether an agent is authorized to exist. It becomes:

Is this agent authorized?

Is this action permitted?

Is this behavior expected for this agent, pursuing this goal, on behalf of this user?

When the Actor Itself Is Dynamic

Agent behavior can also be temporary and difficult to attribute after the fact.

One CISO distinguished relatively durable agents—which can be tied to a user or service principal—from ephemeral sub-agents that may be created for a single task and exist only briefly.

That organization could still observe underlying network activity. But the CISO drew a distinction between seeing traffic and being able to attribute that activity specifically to an agent, a problem that becomes harder when the actor is short-lived or buried several layers deep in another process.

This brings the report's earlier maturity questions into sharper focus. For an AI agent, security teams may need to understand not just whether behavior deviated from history, but whether the deviation is legitimate given the task, permissions, delegated authority, and actions the agent is capable of taking.

Future State: What Happens When an Agent Creates Another Agent?

The scenario is illustrative rather than evidence that this is already a widespread enterprise threat. But it exposes the behavioral challenge: if an agent is expected to exist once within a defined boundary, the appearance of additional instances could itself become a meaningful behavioral anomaly, **provided the organization has defined that expected state in the first place.**

> **How are you gonna go audit an AI activity that lasted for 30 seconds on a virtual machine?"**
>
> — CISO, B2B software organization

One CISO raised a more advanced version of the same problem:

> **What is to stop that agent from cloning itself? And then telling its clone to go do stuff on its behalf?"**

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Building the Controls While Building the Agents

The CISO interviews suggest that organizations are approaching this from very different starting points.

One CISO said the organization had not yet broadly deployed its own AI agents because it was still building both the agents and **"that whole ecosystem for securely managing the AI."** Behavioral monitoring was being developed as part of that process rather than bolted on afterward.

That approach points toward the underlying challenge the survey captures. AI agent security is not waiting for the technology to mature before becoming a budget priority. **Respondents are already planning to prioritize AI agent security investment while the behavioral model is still being worked out.**

That makes behavioral understanding even more urgent.

The next frontier of behavioral security is not simply monitoring more identities. It is defining normal for actors whose legitimate behavior may itself be dynamic, not predictable.

The actor changes. The requirement does not: understand normal well enough to recognize meaningful deviation.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

Conclusion: What Comes After Behavioral Adoption

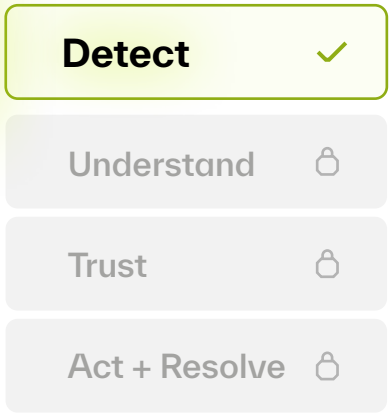
Behavioral security is already firmly established across enterprise security programs. The harder work now is making behavioral understanding precise, trustworthy, and actionable enough to operate at enterprise scale.

That changes what maturity looks like. It is not simply the ability to flag an unusual event. It is the ability to establish enough context around an identity and its activity to determine **what changed, why it matters, and what should happen next.** The research consistently locates the remaining friction in those later stages: from defining normal and reducing noise to explaining anomalies and turning them into confident action.

Detection creates a signal. Behavioral maturity turns that signal into a decision.

Security teams are far past establishing expectations only for people. Machine identities, service accounts, integrations, and increasingly autonomous AI agents are accessing systems, handling information, and taking action. The same fundamental question follows each of them: **what should normal behavior look like here?**

The goal cannot be to treat every deviation as malicious, nor to assume that authorization makes every action safe. The challenge is building enough behavioral understanding to recognize when activity has meaningfully departed from what the identity, system, or agent should be doing.



The future of behavioral security will be defined not by how much activity organizations can observe, but by how well they can baseline normal activity—and act when it changes.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED

About Abnormal AI

Abnormal AI stops cybercrime with AI, protecting more than 5,000 organizations, including 30% of the Fortune 500. The Abnormal Behavioral Security Platform uses proprietary behavioral AI models to understand normal activity across enterprise identities and detect anomalous behavior associated with sophisticated cyber threats.
Learn more at abnormal.ai.

Move From Anomaly Detection to Behavioral Understanding

See how Abnormal uses Behavioral AI to establish what normal looks like across people, machine identities, and autonomous AI agents, distinguish meaningful risk from noise, and turn anomalies into confident action.

Request a Demo ›

Percentages are rounded to the nearest whole number in summary figures. Precise values appear in the body text and charts. Totals may not sum to 100% due to rounding, and totals for multiple-response questions may exceed 100%.

The 2026 State of Behavioral Security report is based on a survey of 525 cybersecurity professionals working at organizations with 1,000 or more employees across six countries. Approximately 85% of respondents held CISO-level or security operations leadership roles. The research also included four in-depth interviews with CISOs. The study was fielded in 2026 by AimPoint Group and commissioned by Abnormal AI.

The findings reflect respondents' reported practices, perceptions, and investment intentions. They should not be interpreted as independently verified measurements of organizational security performance or as representative of all organizations worldwide.